

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

**Patent Application**

Applicant(s): Arun Kwangil Iyengar  
Docket No.: YOR920010663US1  
Serial No.: 10/629,284  
Filing Date: July 29, 2003  
Group: 2188  
Examiner: Duc T. Doan

Title: Methods and Systems for Managing Persistent  
Storage of Small Data Objects

---

**APPEAL BRIEF**

Assistant Commissioner for Patents  
Washington, D.C. 20231

Sir:

Applicants (hereinafter referred to as "Appellants") hereby appeal the final rejection of claims 1-28 of the above-identified application.

**REAL PARTY IN INTEREST**

The present application is assigned to International Business Machines Corporation, as evidenced by an assignment recorded July 29, 2003 in the U.S. Patent and Trademark Office at Reel 14329, Frame 799. The assignee, International Business Machines Corporation, is the real party in interest.

**RELATED APPEALS AND INTERFERENCES**

Applicants are not aware of any related appeals or interferences.

**STATUS OF CLAIMS**

Claims 1-28 stand finally rejected under 35 U.S.C. §103(a). Claims 1-28 are appealed.

STATUS OF AMENDMENTS

There has been no amendment filed subsequent to the final rejection.

SUMMARY OF CLAIMED SUBJECT MATTER

As the background section of the present application explains, a key problem in data storage management is how to efficiently manage objects which are relatively small but need to be stored persistently, such as in disk storage rather than in main memory only. One of the key issues for storing such objects is that the minimum unit of transfer, known as a sector, is often much larger than the minimum object size. By way of example, a sector size may be more than a factor of two larger than some objects. This means that the disk storage system will transfer sectors between disk and main memory. When an object is much smaller than a sector, the disk storage system does not have a method for transferring just the object. The disk storage system has to transfer the entire sector (or sectors if the object is on a sector boundary) containing the object in order to access just the object (Specification, page 1, lines 8-20).

Principles of the present invention provide techniques that efficiently manage small objects in storage transfer units, such as sectors. In one aspect of the invention, a technique for managing storage of objects of sizes smaller than a storage transfer unit in a computer system comprises the following steps. A plurality of storage transfer units are maintained in a first storage medium organized by a quantity of free space in a storage transfer unit. A cache is maintained in a second storage medium, the cache comprising a copy of at least one of said plurality of storage transfer units. In response to a request to store an object of a size less than a storage transfer unit, the technique searches for a cached storage transfer unit with sufficient free space to store the object. If no such cached storage transfer unit can be found, the technique identifies an uncached storage transfer unit with sufficient free space to store the object and stores a copy of the identified storage transfer unit in the cache. The object is stored in the identified storage transfer unit by modifying at least one data structure in the cache and subsequently writing a cached copy of the storage transfer unit to the first storage medium. The first storage medium may comprise disk storage. The second storage medium may comprise main memory (Specification, page 2, lines 2-18).

As further explained in the summary section of the present application, the step of searching for a cached storage transfer unit further may comprise identifying a cached storage transfer unit with sufficient free space to store the object giving preference to such cached storage transfer units with less free space. The step of searching for a cached storage transfer unit may further comprise identifying a cached storage transfer unit with a least amount of free space sufficient to store the object. Further, the step of identifying an uncached storage transfer unit may further comprise identifying an uncached storage transfer unit with sufficient free space giving preference to storage transfer units which minimize fragmentation. The step of identifying an uncached storage transfer unit may still further comprise giving preference to storage transfer units with more free space. The step of identifying an uncached storage transfer unit may yet further comprise identifying a storage transfer unit with a most free space (Specification, page 2, line 19, through page 3, line 2).

Principles of the present invention may further comprise the following steps. At least one list of storage transfer units is maintained. At least one tail pointer to a plurality of contiguous unallocated storage transfer units is maintained. Then, the step of identifying an uncached storage transfer unit further may comprise the steps of searching for an uncached storage transfer unit on the at least one list with sufficient space, and if such an uncached storage transfer unit can not be located, identifying an unallocated storage transfer unit from the at least one tail pointer. The inventive technique may also comprise the step of maintaining a plurality of lists of storage transfer units organized by a quantity of free space in a storage transfer unit. Also, the inventive technique may comprise the step of maintaining at least one tail pointer to a plurality of contiguous unallocated storage transfer units (Specification, page 3, lines 3-13).

Further, a cached copy of a storage transfer unit may be written to the first storage medium in response to at least one of: (i) an object in the storage transfer unit being updated; (ii) a number of changed bytes in the storage transfer unit exceeding a threshold; (iii) a number of changed objects in the storage transfer unit exceeding a threshold; and (iv) the cached copy being about to be purged from the cache. The cached copy may be about to be purged from the cache as a result of at least one of a cache replacement policy and the computer system being about to go down. An application

program may write at least one storage transfer unit to disk in a transactional manner. A storage transfer unit may comprise a sector (Specification, page 3, lines 14-22).

In a second aspect of the invention, a technique for maintaining a plurality of objects in a storage transfer unit comprises the following steps. An object position in the storage transfer unit is identified by an object offset in the storage transfer unit. In response to a request to one of access and update a storage transfer unit, the storage transfer unit is copied so that different objects are copied into different buffers. At least one update to at least one object in the storage transfer unit is performed by modifying at least one buffer. After the at least one update has occurred, the storage transfer unit is updated from the at least one buffer (Specification, page 3, line 23, through page 4, line 3).

Principles of the present invention may further comprise the following steps. When an object which does not have a highest offset is deleted, the offset is added to a list. An allocation request is satisfied by using an offset from the list. If an offset is not reused by the time the storage transfer unit is updated from the at least one buffer, a placeholder is stored on the storage transfer unit indicating the object has been deleted. Still further, the technique may comprise, in response to an object update which would cause a storage transfer unit to overflow, moving the object to a new storage transfer unit and storing a forwarding pointer in the previous storage transfer unit. The storage transfer unit may be maintained on disk and the at least one buffer may be maintained in main memory. The step of updating the storage transfer unit from the at least one buffer may further comprise copying a plurality of objects from buffers to the storage transfer unit in a contiguous area so that free space in the storage transfer unit is contiguous. The storage transfer unit may comprise a sector (Specification, page 4, lines 4-16).

Further, the inventive technique may comprise maintaining a number of free bytes in the storage transfer unit. The step of performing at least one update may further comprise using the number of free bytes in the storage transfer unit to prevent overflow. Also, the copy of at least one of the plurality of storage transfer units included in the cache may be one of a partial copy and an inexact copy (Specification, page 4, lines 17-21).

With respect to an illustrative embodiment, FIG. 1 illustrates a small object disk manager and a storage environment in which it may operate, according to an embodiment of the invention. Disk storage 101 includes a plurality of sectors. A sector is the smallest unit that can be transferred from disk 101 or stored on the disk at a time. Main memory storage 102 includes a cache 103. Cache 103 is used to cache the contents of recently accessed sectors for faster access. Main memory storage 102 can generally be accessed with considerably less overhead than disk storage 101. Therefore, performance advantages are incurred by caching the contents of frequently accessed sectors. As illustrated, a small object disk manager (SODM) 104 is the system which efficiently manages persistent objects of sizes less than a sector. SODM 104 maintains lists of sectors organized by the amount of free storage in a sector. In doing so, SODM 104 keeps track of the amount of free storage within a sector. FIG. 2 shows one method by which SODM 104 manages sectors (Specification, page 6, line 13, through page 7, line 1).

FIGS. 3 through 11 then show further illustrative details of the how SODM 104 manages persistent objects of sizes smaller than a sector (Specification, page 7 through page 12).

#### GROUND OF REJECTION TO BE REVIEWED ON APPEAL

(I) Whether claims 1-3, 9, 11, 12, 14-16 and 19-28 are unpatentable under 35 U.S.C. §103(a) over U.S. Publication No. 2002/0032691 filed in the name of Rabii et al. (hereinafter “Rabii”) in view of U.S. Patent No. 6,915,307 issued to Mattis et al. (hereinafter “Mattis”).

(II) Whether claims 4-8 and 10 are unpatentable under 35 U.S.C. §103(a) over Rabii in view of Mattis in further in view of U.S. Patent No. 6,804,761 issued to Chen et al. (hereinafter “Chen”).

(III) Whether claims 13 and 17 are unpatentable under 35 U.S.C. §103(a) over Rabii in view of Mattis in further in view of U.S. Patent No. 5,802,599 issued to Carbrera et al. (hereinafter “Carbrera”).

(IV) Whether claim 18 is unpatentable under 35 U.S.C. §103(a) over Rabii in view of Mattis in further in view of Carbrera in still further view of U.S. Publication No. 2004/0172507 filed in the name of Garthwaite (hereinafter “Garthwaite”).

ARGUMENT

Appellants incorporate by reference herein the disclosure of their previous response filed in the present application, namely, the response dated November 30, 2005.

(I) Whether claims 1-3, 9, 11, 12, 14-16 and 19-28 are unpatentable under 35 U.S.C. §103(a) over U.S. Publication No. 2002/0032691 filed in the name of Rabii et al. (hereinafter “Rabii”) in view of U.S. Patent No. 6,915,307 issued to Mattis et al. (hereinafter “Mattis”).

Regarding the §103(a) rejections, Applicant asserts that the various references, alone or in combination, fail to teach or suggest all of the limitations of the claims 1-28, as will be explained below. Furthermore, with regard to the combinations of the various references, Applicant asserts that such combinations are improper, as will be explained below.

The Examiner cites Rabii in combination with Mattis in rejecting independent claims 1, 16 and 25-28. More particularly, the Examiner cites portions of Rabii as disclosing certain limitations of the independent claims, and cites portions of Mattis as disclosing certain other limitations of the independent claims. Below Applicant explains how such portions of Rabii and Mattis fail to teach or suggest what the Examiner contends that they teach or suggest. While Applicant may refer from time to time to each reference alone in describing its deficiencies, it is to be understood that such arguments are intended to point out the overall deficiency of the cited combination.

The Rabii patent application is directed to methods for maintaining directory structures on disk. Rabii does not disclose how to allocate space efficiently for small objects as independent claim 1 recites. Claim 1 explicitly recites that it is a method of managing storage of “objects of sizes smaller than a storage transfer unit.”

As pointed out in the background section of the present application and repeated above, a key problem in data storage management is how to efficiently manage objects which are relatively small but need to be stored persistently, such as in disk storage rather than in main memory only, wherein one of the key issues for storing such objects is that the minimum unit of transfer (also referred to as a storage transfer unit), such as a sector, is often much larger than the minimum object size. As explained, a sector size may be more than a factor of two larger than some objects. This

means that the disk storage system will transfer sectors between disk and main memory. When an object is much smaller than a sector, the disk storage system does not have a method for transferring just the object. The disk storage system has to transfer the entire sector (or sectors if the object is on a sector boundary) containing the object in order to access just the object.

Thus, the steps of the claimed invention are directed toward managing storage of “objects of sizes smaller than a storage transfer unit.” Rabii does not introduce any notion of a storage transfer unit or objects of sizes smaller than a storage transfer unit. The final Office Action fails to address this deficiency other than to state (at page 11) that “transferring a fix [sic] amount, for example a cache line between the disk and a cache met the . . . definition.”

Regarding the contention in the final Office Action that Rabii discloses the first step of claim 1 (i.e., maintaining a plurality of storage transfer units in a first storage medium organized by a quantity of free space in a storage transfer unit), Applicant strongly disagrees. First, data object partitions of Rabii are not the same as storage transfer units of the claimed invention. There is no support given in the final Office Action or in Rabii to conclude otherwise. Secondly, Rabii does not organize such data object partitions by a quantity of free space, as the claimed invention recites with respect to storage transfer units.

Furthermore, Rabii fails to disclose the second step of claim 1 (i.e., maintaining in a second storage medium a cache comprising a copy of at least one of said plurality of storage transfer units) since there is no notion of storage transfer units maintained in the data buffers of FIG. 10 of Rabii.

Regarding the contention that Mattis discloses some steps of claim 1, Applicant strongly disagrees. Mattis discloses a method for managing an object cache, i.e., the data objects are stored directly in the cache. The claimed invention recites storing an object of a size less than a storage transfer unit in a storage transfer unit of the cache (which is maintained in the second storage medium) and then writing a cached copy of the storage transfer unit to the first storage medium. This extra level of indirection in the claimed invention further distinguishes it from Mattis.

Still further, the combination of Rabii and Mattis is improper. These two cited references are solving completely different problems. Rabii is directed to the problems associated with maintaining directory structures on a disk, while Mattis is directed to the problems associated with

managing an object cache. Thus, it is not clear how or why one would combine the two disparate references.

The Federal Circuit has stated that when patentability turns on the question of obviousness, the obviousness determination “must be based on objective evidence of record” and that “this precedent has been reinforced in myriad decisions, and cannot be dispensed with.” In re Lee, 277 F.3d 1338, 1343 (Fed. Cir. 2002). Moreover, the Federal Circuit has stated that “conclusory statements” by an examiner fail to adequately address the factual question of motivation, which is material to patentability and cannot be resolved “on subjective belief and unknown authority.” Id. at 1343-1344.

In the final Office Action at page 4, the Examiner provides the following statement to prove motivation to combine Rabii and Mattis, with emphasis supplied: “[i]t would have been obvious to one of ordinary skill in the art at the time of the invention to include the steps for storing an object as suggested by Mattis in Rabii’s system to assure the integrity of information objects . . . .”

Applicant submits that this statement is based on the type of “subjective belief and unknown authority” that the Federal Circuit has indicated provides insufficient support for an obviousness rejection. More specifically, the Examiner fails to identify any objective evidence of record which supports the proposed combination. The Examiner cites Mattis at column 21, lines 1-15, which states that “the foregoing sequence of steps is ordered in a way that ensures the integrity of information objects that are written to the cache.” However, since Rabii has nothing to do with the problems associated with managing an object cache, but rather is directed to the problems associated with maintaining directory structures on a disk, it is unclear how the cited portion of Rabii provides proper motivation to combine Rabii and Mattis. The final Office Action gives no further guidance on this important point.

For at least these reasons, Applicant asserts that independent claim 1 and independent claims 25 and 27 (which contain similar limitations as independent claim 1) are patentable over the Rabii/Mattis combination. Furthermore, Applicant asserts that the claims which depend from claim 1 are patentable over the Rabii/Mattis combination not only for the reasons given above with respect



to claim 1, but also because such dependent claims recite patentable subject matter in their own right, as will be set out below.

Regarding independent claim 16, Applicant asserts that neither Rabii nor Mattis, alone or in combination, teach or suggest all of the limitations of the claimed invention. Applicant notes that the final Office Action corrects the wrong reference cited in the non-final Office Action (Rabii when the non-final Office Action apparently meant Mattis) in the one paragraph argument against claim 16. However, the rationale for rejecting the claim in the final Office Action is still deficient since various limitations in claim 16 are not disclosed by the cited combination. By way of example only, there is nothing the same or analogous in the cited combination to the claimed features of “storage transfer units” and “object offsets,” and how they relate to one another. Also, Rabii is silent to copying a storage transfer unit, in response to a request to one of access and update a storage transfer unit, so that different objects are copied into different buffers.

Further, the Rabii/Mattis combination is improper, as explained above.

For at least these reasons, Applicant asserts that independent claim 16 and independent claims 26 and 28 (which include similar limitations as independent claim 16) are patentable over the Rabii/Mattis combination. Furthermore, Applicant asserts that the claims which depend from claim 16 are patentable over the Rabii/Mattis combination not only for the reasons given above with respect to claim 16, but also because such dependent claims recite patentable subject matter in their own right, as will be set out below.

Regarding claims 2 and 3, whether or not Rabii refers to disk storage and main memory is not relevant, since Rabii fails to refer to specific steps for managing storage of objects of sizes smaller than a storage transfer unit in the context of disk storage and main memory, as in the claimed invention.

Regarding claims 9 and 11, any disclosed data partitions or pointers in Rabii fail to function in the specific ways that claims 9 and 11 respectively recite.

Regarding claim 12, it is not clear how the fact that Rabii describes that “objects remain in cache until there is a shortage of in-memory-objects or data buffers . . . [wherein] the object is purged from cache,” teaches or suggests that a cached copy of a storage transfer unit (from the

second storage medium) is written to the first storage medium in response to at least one of the criterion recited in claim 12.

Regarding claim 14, it is unclear how a “large write back” to the disk in Rabii during a system shutdown supports a rejection of the claimed features of purging the cached copy from the cache as a result of at least one of a cache replacement policy and the computer system being about to go down.

Regarding claims 19 and 20-24, since Rabii does not teach or suggest the claimed features of “storage transfer units” and “object offsets” or the step of copying a storage transfer unit, in response to a request to one of access and update a storage transfer unit, so that different objects are copied into different buffers, it is not clear how the features specifically defined in claims 19 and 20-24 could be taught or suggested by the cited portions of Rabii.

(II) Whether claims 4-8 and 10 are unpatentable under 35 U.S.C. §103(a) over Rabii in view of Mattis in further in view of U.S. Patent No. 6,804,761 issued to Chen et al. (hereinafter “Chen”).

Applicant asserts that the claims 4-8 and 10, which depend from claim 1, are patentable over the Rabii/Mattis/Chen combination not only for the reasons given above with respect to claim 1, but also because such dependent claims recite patentable subject matter in their own right.

Whether or not Chen describes a “chunk manager capable of searching for the next highest size standard memory blocks to fit an application request,” or “searching through a range of blocks with different sizes . . . thereby reducing fragmentation,” or “memory pools with lists of standard blocks,” is not relevant to the rejection of the claimed features in claims 4-8 and 10 due at least to the failure of any of the references in the cited combination to recognize an object of a size less than a storage transfer unit.

It is also asserted that the motivation set forth by the Examiner to combine Chen with Rabii and Mattis is insufficient under In re Lee decision (cited above).

(III) Whether claims 13 and 17 are unpatentable under 35 U.S.C. §103(a) over Rabii in view of Mattis in further in view of U.S. Patent No. 5,802,599 issued to Cabrera et al. (hereinafter “Cabrera”).

Applicant asserts that the claims 13 and 17, which depend from claims 1 and 16 respectively, are patentable over the Rabii/Mattis/Cabrera combination not only for the reasons given above with respect to claims 1 and 16, but also because such dependent claims recite patentable subject matter in their own right.

Cabrera does not disclose the notion of an “object offset.” The portion of Cabrera disclosing reuse of a memory buffer fails to teach or suggest such feature.

It is also asserted that the motivation set forth by the Examiner to combine Cabrera with Rabii and Mattis is insufficient under In re Lee decision (cited above).

(IV) Whether claim 18 is unpatentable under 35 U.S.C. §103(a) over Rabii in view of Mattis in further in view of Cabrera in still further view of U.S. Publication No. 2004/0172507 filed in the name of Garthwaite (hereinafter “Garthwaite”).

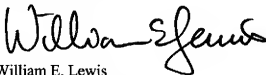
Applicant asserts that the claim 18, which depends from claim 16, is patentable over the Rabii/Mattis/Cabrera/Garthwaite combination not only for the reasons given above with respect to claim 1, but also because such dependent claim recites patentable subject matter in their own right.

First, Cabrera does not perform the steps that the Examiner suggests it does on page 10 of the final Office Action. Further, Garthwaite does not disclose the notion of an “storing a forwarding pointer in the previous storage transfer unit.”

It is also asserted that the motivation set forth by the Examiner to combine Garthwaite with Rabii, Mattis and Cabrera is insufficient under In re Lee decision (cited above).

In view of the above, Applicants believe that claims 1-28 are in condition for allowance, and respectfully request withdrawal of the various §103(a) rejections.

Respectfully submitted,

A handwritten signature in black ink, appearing to read "William E. Lewis". The signature is fluid and cursive, with the first name "William" being more prominent and the last name "Lewis" following in a similar style.

Date: October 10, 2006

William E. Lewis  
Attorney for Applicant(s)  
Reg. No. 39,274  
Ryan, Mason & Lewis, LLP  
90 Forest Avenue  
Locust Valley, NY 11560  
(516) 759-2946

APPENDIX

1. A method of managing storage of objects of sizes smaller than a storage transfer unit in a computer system, comprising the steps of:

maintaining a plurality of storage transfer units in a first storage medium organized by a quantity of free space in a storage transfer unit;

maintaining in a second storage medium a cache comprising a copy of at least one of said plurality of storage transfer units;

in response to a request to store an object of a size less than a storage transfer unit:

searching for a cached storage transfer unit with sufficient free space to store the object;

if no such cached storage transfer unit can be found, identifying an uncached storage transfer unit with sufficient free space to store the object and storing a copy of the identified storage transfer unit in the cache; and

storing the object in the identified storage transfer unit by modifying at least one data structure in the cache and subsequently writing a cached copy of the storage transfer unit to the first storage medium.

2. The method of claim 1, wherein the first storage medium comprises disk storage.

3. The method of claim 1, wherein the second storage medium comprises main memory.

4. The method of claim 1, wherein the step of searching for a cached storage transfer unit further comprises identifying a cached storage transfer unit with sufficient free space to store the object giving preference to such cached storage transfer units with less free space.

5. The method of claim 4, wherein the step of searching for a cached storage transfer unit further comprises identifying a cached storage transfer unit with a least amount of free space sufficient to store the object.

6. The method of claim 1, wherein the step of identifying an uncached storage transfer unit further comprises identifying an uncached storage transfer unit with sufficient free space giving preference to storage transfer units which minimize fragmentation.

7. The method of claim 1, wherein the step of identifying an uncached storage transfer unit further comprises giving preference to storage transfer units with more free space.

8. The method of claim 7, wherein the step of identifying an uncached storage transfer unit further comprises identifying a storage transfer unit with a most free space.

9. The method of claim 1, further comprising the steps of:  
maintaining at least one list of storage transfer units;  
maintaining at least one tail pointer to a plurality of contiguous unallocated storage transfer units;

wherein the step of identifying an uncached storage transfer unit further comprises the steps of searching for an uncached storage transfer unit on the at least one list with sufficient space, and if such an uncached storage transfer unit can not be located, identifying an unallocated storage transfer unit from the at least one tail pointer.

10. The method of claim 1, further comprising the step of maintaining a plurality of lists of storage transfer units organized by a quantity of free space in a storage transfer unit.

11. The method of claim 1, further comprising the step of maintaining at least one tail pointer to a plurality of contiguous unallocated storage transfer units.

12. The method of claim 1, wherein a cached copy of a storage transfer unit is written to the first storage medium in response to at least one of: (i) an object in the storage transfer unit being updated; (ii) a number of changed bytes in the storage transfer unit exceeding a threshold; (iii) a

number of changed objects in the storage transfer unit exceeding a threshold; and (iv) the cached copy being about to be purged from the cache.

13. The method of claim 1, wherein an application program writes at least one storage transfer unit to disk in a transactional manner.

14. The method of claim 12, wherein the cached copy is about to be purged from the cache as a result of at least one of a cache replacement policy and the computer system being about to go down.

15. The method of claim 1, wherein a storage transfer unit comprises a sector.

16. A method of maintaining a plurality of objects in a storage transfer unit, comprising the steps of:

identifying an object position in the storage transfer unit by an object offset in the storage transfer unit;

in response to a request to one of access and update a storage transfer unit, copying the storage transfer unit so that different objects are copied into different buffers;

performing at least one update to at least one object in the storage transfer unit by modifying at least one buffer; and

after the at least one update has occurred, updating the storage transfer unit from the at least one buffer.

17. The method of claim 16, further comprising the steps of:

when an object which does not have a highest offset is deleted, adding the offset to a list;

satisfying an allocation request by using an offset from the list; and

if an offset is not reused by the time the storage transfer unit is updated from the at least one buffer, storing a placeholder on the storage transfer unit indicating the object has been deleted.

18. The method of claim 16, further comprising the step of:

in response to an object update which would cause a storage transfer unit to overflow, moving the object to a new storage transfer unit and storing a forwarding pointer in the previous storage transfer unit.

19. The method of claim 16, wherein the storage transfer unit is maintained on disk and the at least one buffer is maintained in main memory.

20. The method of claim 16, wherein the step of updating the storage transfer unit from the at least one buffer further comprises copying a plurality of objects from buffers to the storage transfer unit in a contiguous area so that free space in the storage transfer unit is contiguous.

21. The method of claim 16, wherein the storage transfer unit comprises a sector.

22. The method of claim 16, further comprising the step of maintaining a number of free bytes in the storage transfer unit.

23. The method of claim 22, wherein the step of performing at least one update further comprises using the number of free bytes in the storage transfer unit to prevent overflow.

24. The method of claim 16, wherein the copy of at least one of the plurality of storage transfer units included in the cache is one of a partial copy and an inexact copy.

25. Apparatus for managing storage of objects of sizes smaller than a storage transfer unit in a computer system, comprising:

at least one processor operative to: (i) maintain a plurality of storage transfer units in a first storage medium organized by a quantity of free space in a storage transfer unit; (ii) maintain in a



second storage medium a cache comprising a copy of at least one of said plurality of storage transfer units; (iii) in response to a request to store an object of a size less than a storage transfer unit: searching for a cached storage transfer unit with sufficient free space to store the object; if no such cached storage transfer unit can be found, identifying an uncached storage transfer unit with sufficient free space to store the object and storing a copy of the identified storage transfer unit in the cache; and storing the object in the identified storage transfer unit by modifying at least one data structure in the cache and subsequently writing a cached copy of the storage transfer unit to the first storage medium.

26. Apparatus for maintaining a plurality of objects in a storage transfer unit, comprising: at least one processor operative to: (i) identify an object position in the storage transfer unit by an object offset in the storage transfer unit; (ii) in response to a request to one of access and update a storage transfer unit, copy the storage transfer unit so that different objects are copied into different buffers; (iii) perform at least one update to at least one object in the storage transfer unit by modifying at least one buffer; and (iv) after the at least one update has occurred, update the storage transfer unit from the at least one buffer.

27. An article of manufacture for managing storage of objects of sizes smaller than a storage transfer unit in a computer system, comprising a machine readable medium containing one or more programs which when executed implement the steps of:

maintaining a plurality of storage transfer units in a first storage medium organized by a quantity of free space in a storage transfer unit;

maintaining in a second storage medium a cache comprising a copy of at least one of said plurality of storage transfer units;

in response to a request to store an object of a size less than a storage transfer unit:

searching for a cached storage transfer unit with sufficient free space to store the object;

if no such cached storage transfer unit can be found, identifying an uncached storage transfer unit with sufficient free space to store the object and storing a copy of the identified storage transfer unit in the cache; and

storing the object in the identified storage transfer unit by modifying at least one data structure in the cache and subsequently writing a cached copy of the storage transfer unit to the first storage medium.

28. An article of manufacture for maintaining a plurality of objects in a storage transfer unit, comprising a machine readable medium containing one or more programs which when executed implement the steps of:

identifying an object position in the storage transfer unit by an object offset in the storage transfer unit;

in response to a request to one of access and update a storage transfer unit, copying the storage transfer unit so that different objects are copied into different buffers;

performing at least one update to at least one object in the storage transfer unit by modifying at least one buffer; and

after the at least one update has occurred, updating the storage transfer unit from the at least one buffer.

EVIDENCE APPENDIX

None.

RELATED PROCEEDINGS APPENDIX

None.